

**ROBUST SYSTEM BUS RECOVERY****BACKGROUND OF THE INVENTION**

5

**Cross-reference to Related Applications**

10

The present application is related to co-pending U.S. Patent Application Serial No. 09/436,898, "Multi-Node Data Processing System Having a Non-Hierarchical Interconnect Architecture" filed on November 9, 1999, assigned to the assignee of the present application and incorporated herein by reference.

15

**1. Technical Field:**

20

The present invention relates in general to data processing and, in particular, to an interconnect of a data processing system. Still more particularly, the present invention relates to data processing systems of processing nodes having recovery methods. The nodes can be arranged to operate either in a multi-node data processing system having a non-hierarchical interconnect architecture topology, or on a different topology, such as over a common hierarchical bus.

25

**2. Description of the Related Art:**

30

It is well-known in the computer arts that greater computer system performance can be achieved by harnessing the processing power of multiple individual processors in tandem. Multi-processor (MP) computer systems can be designed with a number of different architectures, of which various ones may be better suited for particular applications depending upon the intended design point, the system's performance requirements, and the software environment of each application. Known architectures include, for example, the symmetric multiprocessor (SMP) and non-uniform memory access (NUMA) architectures. Until the present invention, it has generally been assumed that greater scalability and hence greater performance is obtained by designing more hierarchical computer systems, that is,

computer systems having more layers of interconnects and fewer processor connections per interconnect.

5 The present invention recognizes, however, that such hierarchical computer systems incur extremely high communication latency for the percentage of data requests and other transactions that must be communicated between processors coupled to different interconnects. For example, even for the relatively simple case of an 8-way SMP system in which four processors present in each of two nodes are coupled by an upper level bus and the two nodes are themselves coupled by a lower level bus, communication of a data request between processors in different nodes will incur bus acquisition and other transaction-related latency at each of three buses. 10 Because such latencies are only compounded by increasing the depth of the interconnect hierarchy, the present invention recognizes that it would be desirable and advantageous to provide an improved data processing system architecture having reduced latency for transaction between physically remote processors.

15 The present invention additionally recognizes that from time to time errors occur in processing in data processing systems even those operating in high speed, high frequency bandwidth topologies. Normally, it would be expected that a system processing error in such topologies would cause an overall system failure, requiring a time consuming effort for system recovery at the high frequency. It would thus be desirable to prove a method and system for more robust recovery in high speed, high bandwidth data processing systems. 20

### SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide for recovery in a data processing system in the event of system errors.

It is a further object of the present invention to provide a data processing system and method operating on a high speed topology capable of operating while system errors are being corrected.

It is yet another object of the present invention to provide a dynamic ability to increase system performance in high speed, high performance data processing systems.

The present invention realizes the above and other advantages in a multi-node data processing system having a non-hierarchical interconnect architecture. The foregoing and other objects and advantages of the present invention will be apparent to those skilled in the art, in view of the following detailed description of the preferred embodiment of the present invention, taken in conjunction with the appended claims and the accompanying drawings.

In accordance with the present invention, a data processing system includes a plurality of nodes, which each contain at least one agent, and data storage accessible to agents within the nodes. The nodes are coupled by a high speed, high bandwidth topology as a system topology. Additionally included in the system topology is a recovery bus operating at a lower speed than the high speed topology and connecting the processing nodes together. The agents in the nodes monitor the status of processing in the high speed topology to sense errors. When an error is sensed, communication is transferred to the low speed recovery bus. The topology preferably takes the form of an interconnect including a plurality of address channels to which each agent is coupled and at least one data channel. Each agent can only issue transactions on an associated address channel. However, agents snoop transactions on all of the plurality of address channels.

All objects, features, and advantages of the present invention will become apparent in the following detailed written description.

### BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives, and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

**Figure 1** depicts an illustrative embodiment of a multi-node data processing system having a non-hierarchical interconnect architecture in accordance with the present invention;

**Figure 2** is a more detailed block diagram of a processor embodiment of an agent within the data processing system of **Figure 1**;

**Figure 3** is a more detailed block diagram of the communication logic of the processor in **Figure 2**;

**Figure 4** is a more detailed block diagram of response and flow control logic within the data processing system shown in **Figure 1**;

**Figure 5A** is a timing diagram of an exemplary address transaction in the data processing system illustrated in **Figure 1**;

**Figure 5B** is a timing diagram of an exemplary read-data transaction in the data processing system depicted in **Figure 1**;

**Figure 5C** is a timing diagram of an exemplary write-data transaction in the data processing system illustrated in **Figure 1**;

**Figure 6A** depicts an exemplary format of a request transaction transmitted via one of the address channels of the data processing system shown in **Figure 1**;

**Figure 6B** illustrates an exemplary format of a partial combined response or combined response transmitted via one of the response channels of the data processing system of **Figure 1**;

**Figure 6C** depicts an exemplary format of a data transaction transmitted via the data channel of the data processing system of **Figure 1**;

**Figure 7** illustrates an alternative embodiment of a multi-node data processing system having a non-hierarchical interconnect architecture in accordance with the present invention; and

**Figure 8** is a flow diagram of the operation of a system in a recovery mode according to the present invention.

5

FIG. 7

## DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

With reference now to the figures and in particular with reference to Figure 1, there is depicted an illustrative embodiment of a multi-node data processing system 8 having a non-hierarchical interconnect architecture in accordance with the present invention. As shown, data processing system 8 includes a number of nodes 10a-10k, which are coupled together in a ring configuration by a segmented interconnect 12 having one segment per node 10.

In addition to a segment of interconnect 12, each node 10 of data processing system 8 includes one or more agents that are each coupled to interconnect 12 and are designated A0-An for node 10a, B0-Bn for node 10b, etc. Each node 10 also includes respective response and flow control logic 18 that controls the flow of transactions on interconnect 12 between its node 10 and a neighboring node 10 and generates sideband signals (discussed below) that indicate how agents snooping a request should respond. The number of agents within each node 10 is preferably limited to an interconnect-dependent performance-optimized number (e.g., 8 or 16), with greater system scale being achieved by adding additional nodes 10 to data processing system 8.

According to the present invention, the interconnect 12 of the agents A0-An for node 10a, B0-Bn for node 10b, etc., is configured to operate as a high speed, high bandwidth or wide bus, operating at a frequency of, for example, on the order of 500Mhz or higher. Configuration registers in the agents maintain record of this topology. The interconnect 12 is also selectively configurable on command, as will be set forth, to select and interconnect a designated one of the set of agents, such as A1, B1, etc., or some higher number of the sets of agents, for each one of the nodes 10 on a lower frequency (such as 125 Mhz or 250 Mhz or some other submultiple frequency of the frequency of the wide bus), as part of a narrow bandwidth recovery bus L according to the present invention. The recovery bus L is indicated in phantom in Figure 1. As will be set forth, the number of nodes connected on the low frequency recovery bus L is selectively definable and may include each of the nodes or only

selected ones. The configuration registers previously mentioned maintain record of whether the interconnect 12 is to operate on the high speed, wide bus or the recovery bus based on the state of data processing in the system 8.

Turning now more specifically to the interconnect architecture of data processing system 8, interconnect 12 includes at least one (and in the illustrated embodiment a single) data channel 16 and a plurality of non-blocking address channels 14a-14k that are each associated with a respective one of nodes 10a-10k such that only agents within the associated node 10 can issue requests on an address channel 14. Each of address channels 14 and data channel 16 of interconnect 12 is segmented, as noted above, such that each node 10 contains a segment of each address and data channel, and each address and data channel segment is coupled (by flow control logic) to at least two neighboring segments of the same channel. As indicated by arrows, each channel is also uni-directional, meaning that address and data transactions on interconnect 12 are only propagated between neighboring nodes 10 in the indicated direction. In the illustrated embodiment, each segment of an address channel 14 is implemented as an address bus that conveys 32 address bits in parallel, and each segment of data channel 16 is implemented as a data bus that conveys 16 data bytes in parallel; however, it will be appreciated that individual segments of interconnect 12 can alternatively be implemented with switch-based or hybrid interconnects and that other embodiments of the present invention may implement different channel widths.

In conjunction with interconnect 12, data processing system 8 implements three sideband channels -- a partial combined response channel 24, a combined response channel 26, and a cancel channel 27 -- to respectively communicate partial combined responses, combined responses, and a cancel (or stomp) signal. As utilized herein, a partial combined response (or PCR) is defined as a cumulative response to a request of all agents within fewer than all nodes, and a combined response (or CR) is defined as a cumulative response to a request by all agents in all nodes. As discussed further below, agents are able to determine by reference to the PCR, CR, and cancel

signal associated with a request snooped on an address channel 14 whether or not to service the request.

Referring now to Figure 2, there is depicted a block diagram of a processor 28 that can be utilized to implement any agent within data processing system 8.

Although hereafter it is assumed that each agent within data processing system 8 is a processor, it should be understood that an agent can be any device capable of supporting the communication protocol described herein.

As shown in Figure 2, processor 28 includes processing logic 30 for processing instructions and data, communication logic 34, which implements a communication protocol that governs communication on interconnect 12, and a cache hierarchy 32 that provides local, low latency storage for instructions and data. In addition to cache hierarchy 32, which may include, for example, level one (L1) and level two (L2) caches, the local storage of each processor 28 may include an associated off-chip level three (L3) cache 20 and local memory 22, as shown in Figure 1. Instructions and data are preferably distributed among local memories 22 such that the aggregate of the contents of all local memories 22 forms a shared "main memory" that is accessible to any agent within data processing system 8. Hereinafter, the local memory 22 containing a storage location associated with a particular address is said to be the home local memory for that address, and the agent interposed between the home local memory and interconnect 12 is said to be the home agent for that address. As shown in Figure 2, each home agent has a memory map 36 accessible to cache hierarchy 32 and communication logic 34 that indicates only what memory addresses are contained in the attached local memory 22.

With reference now to Figure 3, there is illustrated a more detailed block diagram representation of an illustrative embodiment of communication logic 34 of Figure 2. As illustrated, communication logic 34 includes master circuitry comprising master control logic 40, a master address sequencer 42 for sourcing request (address) transactions on an address channel 14, and a master data sequencer 44 for sourcing data transactions on data channel 16. Importantly, to ensure that each of address



channels 14 is non-blocking, the master address sequencer 42 of each agent within a given node 10 is connected to only the address channel 14 associated with its node 10. Thus, for example, the master address sequencer 42 of each of agents A0-An is connected to only address channel 14a, the master address sequencer 42 of each of agents B0-Bn is connected to only address channel 14b, and the master address sequencer 42 of each of agents K0-Kn is connected to only address channel 14k. To fairly allocate utilization of address channels 14 and ensure that local agents do not issue conflicting address transactions, some arbitration mechanism (e.g., round robin or time slice) should be utilized to arbitrate between agents within the same node 10.

By contrast, the master data sequencers 44 of all agents within data processing system 8 are connected to data channel 16. Although a large number of agents may be connected to data channel 16, in operation data channel 16 is also non-blocking since the types of data transactions that may be conveyed by data channel 16, which predominantly contain (1) modified data sourced from an agent other than the home agent, (2) data sourced from the home agent, and (3) modified data written back to the home local memory 22, are statistically infrequent for applications in which the distribution of memory among local memories 22 and the distribution of processes among the agents is optimized. Of course, in implementations including only a single data channel 16, some arbitration mechanism (e.g., round robin or time slice) should be utilized to arbitrate between agents within the same node 10 to ensure that local agents do not issue conflicting data transactions.

Communication logic 34 also includes snoop circuitry comprising a snoop address and response sequencer 52 coupled to each address channel 14 and to sideband response channels 24 and 26, a snoop data sequencer 54 coupled to data channel 16, and snoop control logic 50 connected to snoop address and response sequencer 52 and to snoop data sequencer 54. In response to receipt of a request transaction by snoop address and response sequencer 52 or a data transaction by snoop data sequencer 54, the transaction is passed to snoop control logic 50. Snoop control logic 50 processes the transaction in accordance with the

implemented communication protocol and, if a request transaction, provides a snoop response and possibly a cancel signal to its node's response and flow control logic 18. Depending upon the type of transaction received, snooper control logic 50 may initiate an update to a directory or data array of cache hierarchy 32, a write to the local memory 22, or some other action. Snooper control logic 50 performs such processing of request and data transactions from a set of request queues 56 and data queues 58, respectively.

Referring now to Figure 4, there is depicted a more detailed block diagram of an exemplary embodiment of response and flow control logic 18. As illustrated, response and flow control logic 18 includes response logic 60, which combines snoop responses from local agents and possibly a PCR from a neighboring node 10 to produce a cumulative PCR indicative of the partial combined response for all nodes that have received the associated transaction. For example, if agent A0 of node 10a masters a request on address channel 14a, agents A1-An provide snoop responses that are combined by response and flow control logic 18a to produce a  $PCR_A$  that is provided on PCR bus 24. When the request is snooped by agents B0-Bn, agents B0-Bn similarly provide snoop responses, which are combined with  $PCR_A$  of node 10a by response and flow control logic 18b to produce a cumulative  $PCR_{A+B}$ . This process continues until a complete combined response is obtained (i.e.,  $PCR_{A+B+...+K} = CR$ ). Once the CR is obtained, the CR is made visible to all nodes via CR channel 26. Depending upon the desired implementation, the CR for a request can be provided on CR channel 26 by the response and flow control logic 18 of either the last node 10 receiving the request or the master node 10 containing the master agent. It is presently preferable, both in terms of complexity and resource utilization, for the response logic 60 of the master node 10 to provide the CR for a request, thus permitting agents within the master node 10 to receive the CR prior to agents within any other node 10. This permits the master agent, for example, to retire queues in master control logic 40 which are allocated to the request as soon as possible.

As is further illustrated in Figure 4, response and flow control logic 18 also contains internal or address flow control logic 62, which includes address latches 64 connecting neighboring segments of each of address channels 14a-14k. Address latches 64 are enabled by an enable signal 66, which can be derived from an interconnect clock, for example. Address flow control logic 62 also includes a data latch 72 that connects neighboring segments of data channel 16. As indicated by enable logic including XOR gate 68 and AND gate 70, data latch 72 operates to output a data transaction to the neighboring segment of data channel 16 only if a the data transaction's destination identifier (ID) does not match the unique node ID of the current node 10 (i.e., if the data transaction specifies an intended recipient node 10 other than the current node 10). Thus, data transactions communicated on data channel 16, which can contain either read data or write data, propagate from the source node to the destination node (which may be the same node), utilizing only the segments of data channel 16 within these nodes and any intervening node(s) 10.

Each response and flow control logic 18 further includes cancellation logic 74, which is implemented as an OR gate 76 in the depicted embodiment. Cancellation logic 74 has an output coupled to cancel channel 27 and an input coupled to the cancel signal output of the snooper control logic 50 of each agent within the local node 10. The snooper control logic 50 of an agent asserts its cancel signal if the snooper control logic 50 determines, prior to receiving the PCR from another node 10, that a request issued by an agent within the local node 10 will be serviced by an agent within the local node 10. Depending on the desired implementation, the cancel signal can be asserted by either or both of the master agent that issued the request and the snooping agent that will service the request. In response to the assertion of the cancel signal of any agent within the node 10 containing the master agent, cancellation logic 74 asserts a cancel signal on cancel channel 27, which instructs the snooper control logic 50 of agents in each other node 10 to ignore the request. Thus, the assertion of a cancel signal improves the queue utilization of agents in remote

nodes 10 by preventing the unnecessary allocation of request and data queues 56 and 58.

With reference now to Figure 5A, a timing diagram of an exemplary request transaction in the data processing system of Figure 1 is depicted. The request transaction is initiated by a master agent, for example, agent A0 of node 10a, mastering a read or write request transaction on the address channel 14 associated with its node, in this case address channel 14a. As shown in Figure 6A, the request transaction 80 may contain, for example, a master node ID field 82 indicating the node ID of the master agent, a transaction type (TT) field 84 indicating whether the request transaction is a read (e.g., read-only or read-with-intent-to-modify) or write request, and a request address field 86 specifying the request address. The request transaction propagates sequentially from node 10a to node 10b and eventually to node 10k via address channel 14a. Of course, while the request transaction is propagating through other nodes 10, other request transactions may be made concurrently on address channel 10a or address channels 14b-14k.

As discussed above and as shown in Figure 5A, after the snooper address and response sequencer 52 of each agent snoops the request transaction on address channel 14a, the request transaction is forwarded to snooper control logic 50, which provides to the local response and flow control logic 18 an appropriate snoop response indicating whether that agent can service (or participate in servicing) the request. Possible snoop responses are listed in Table I below in order of descending priority.

TABLE I

Snoop response	Meaning
Retry	Retry transaction
Modified intervention	Agent holds requested line in a modified state in cache from which data can be sourced
Shared intervention	Agent holds requested line in a shared state from which data can be sourced
Shared	Agent holds requested line in a shared state in cache
Home	Agent is home agent of request address
Null	Agent does not hold the requested line in cache and is not the home agent

The snoop responses of only agents A0-Ak are then combined by response and flow control logic 18a into a  $PCR_A$  output on PCR channel 24. As indicated in Figure 6B, a response 90, which may be either a PCR or a CR, includes at least a response field 94 indicating the highest priority snoop response yet received and a snoop node ID field 92 indicating the node ID of the agent providing the highest priority snoop response yet received.

If during a determination of the appropriate snoop response, the snoop control logic 50 of an agent within node 10a determines that it is likely to have the highest priority snoop response of all agents within data processing system 8, for example, Modified Intervention for a read request or Home for a write request, the agent within node 10a asserts its cancel signal to the local cancellation logic 74, which outputs a cancel signal on cancel channel 27. As shown in Figure 5A, the cancel signal is preferably asserted on cancel channel 27 prior to  $PCR_A$ . Thus, each agent within the nodes that subsequently receive the request transaction (i.e., nodes 10b-10k) can cancel the request queue 56 that is allocated within snoop control

logic 50 to provide the snoop response for the request, and no other snoop responses and no PCR or CR will be generated for the request transaction.

Assuming that no agent within the master node 10a asserts its cancel signal to indicate that the request transaction will be serviced locally, agents B0-Bn within neighboring node 10b will provide snoop responses, which are combined together with  $PCR_A$  by response and flow control logic 18b to produce  $PCR_{A+B}$ . The process of accumulating PCRs thereafter continues until response and flow control logic 18k produces  $PCR_{A+B+...+K}$ , which contains the node ID of the agent that will participate in servicing the request transaction and the snoop response of that servicing agent. Thus, for a read request, the final PCR contains the node ID of the agent that will source the requested cache line of data, and for a write request, the final PCR specifies the node ID of the home agent for the requested cache line of data. When  $PCR_{A+B+...+K}$ , which is equivalent to the CR, is received by response logic 60 within node 10a, response logic 60 of node 10a provides the CR to all agents on CR channel 26.

As illustrated in Figures 1 and 3, each agent within data processing system 8 is coupled to and snoops PCRs on PCR channel 24. In contrast to conventional multi-processor systems in which processors only receive CRs, the present invention makes PCRs visible to agents to permit agents that are not likely to service a snooped request to speculatively cancel queues (e.g., request and/or data queues 56 and 58) allocated to the request prior to receipt of the CR for the request. Thus, if an agent provides a lower priority snoop response to a request than is indicated in the PCR, the agent can safely cancel any queues allocated to the request prior to receiving the CR. This early deallocation of queues advantageously increases the effective size of each agent's queues.

With reference now to Figures 5B and 5C, there are respectively illustrated timing diagrams of an exemplary read-data transaction and an exemplary write-data transaction in data processing system 8 of Figure 1. Each of the illustrated data transactions follows a request (address) transaction such as that illustrated in Figure

5A and assumes agent B0 of node 10b participates with agent A0 of node 10a in the data transaction.

Referring first to the read-data transaction shown in Figure 5B, when the CR output on CR channel 26 by response and flow control logic 18a is received by agent B0, agent B0, which responded to the request transaction with a Modified Intervention, Shared Intervention or Home snoop response indicating that agent B0 could source the requested data, sources a data transaction on data channel 16 containing a cache line of data associated with the request address. As illustrated in Figure 6C, in a preferred embodiment a read-data or write-data transaction 100 includes at least a data field 104 and a destination node ID field 102 specifying the node ID of the node 10 containing the intended recipient agent (in this case node 10a). For read-data requests such as that illustrated in Figure 5B, the destination node ID is obtained by the source agent from master node ID field 82 of the request transaction.

The data transaction sourced by agent B0 is then propagated via data channel 16 through each node 10 until node 10a is reached. As indicated in Figure 5B, response and flow control logic 18a of node 10a does not forward the data transaction to node 10b since the destination node ID contained in field 102 of the data transaction matches the node ID of node 10a. Snooper data sequencer 54 of agent A0 finally snoops the data transaction from data channel 16 to complete the data transaction. The cache line of data may thereafter be stored in cache hierarchy 32 and/or supplied to processing logic 30 of agent A0.

Referring now to Figure 5C, a write-data transaction begins when agent A0, the agent that mastered the write request, receives the CR for the write request via CR channel 26. Importantly, the CR contains the node ID of the home agent of the request address (in this case the node ID of node 10b) in snooper node ID field 92, as described above. Agent A0 places this node ID in destination node ID field 102 of a write-data transaction and sources the data transaction on data channel 16. As indicated in Figure 5C, response and flow control logic 18b of node 10b does not forward the data transaction to any subsequent neighboring node 10 since the

destination node ID contained in field 102 of the data transaction matches the node ID of node 10b. Snooper data sequencer 54 of agent B0 finally snoops the data transaction from data channel 16 to complete the data transaction. The data may thereafter be written into local memory 22 of agent B0.

Importantly, the write-data transaction protocol described above, which is characterized by the target agent being identified (e.g., by device ID, bus ID, node ID, etc.) in the combined response to the request (address) portion of the write transaction and the master subsequently outputting the target agent ID in conjunction with the data portion of the write transaction to route or facilitate snooping of the write data, is not limited to multi-node data processing systems or data processing system embodiments having segmented data channels. In fact, this write-data transaction protocol is generally applicable to inter-chip communication in multiprocessor computer systems and inter-processor communication in single chip multiprocessor systems.

With reference now to Figure 7, there is illustrated an alternative embodiment of a multi-node data processing system having a non-hierarchical interconnect architecture in accordance with the present invention. As shown, data processing system 108, like data processing system 8 of Figure 1, includes a number of nodes 10a-10k, which are coupled together in a ring configuration by a segmented interconnect 112 having one segment per node 10. Interconnect 112 includes at least one (and in the illustrated embodiment a single) data channel 16 and a plurality of non-blocking address channels 14a-14n that are each associated with a particular agent (or connection for an agent) in each one of nodes 10a-10k, such that only agents with the corresponding numerical designation can issue requests on an address channel 14. That is, although each agent snoops all address channels 14, only agents A0, B0, ..., K0 can issue requests on address channel 14a, and only agents An, Bn, ..., Kn can issue requests on address channel 14n. Thus, the principal difference between the embodiments depicted in Figures 1 and 7 is the centralization of master agents for a particular address channel 14 within a single node in Figure 1 versus the one-per-



node distribution of master agents for a particular address channel 14 among nodes 10 in Figure 7.

One advantage of the interconnect architecture illustrated in Figure 7 is that master agents need not arbitrate for their associated address channels 14. If the snooper control logic 50 of an agent detects that no address transaction is currently being received on the associated address channel, the master control logic 40 can source an address transaction on its address channel 14 without the possibility of collision with another address transaction.

According to the present invention, the interconnect 12 takes the form of a high speed, high bandwidth or wide bus topology, preferably configured in a loop topology, as illustrated, at frequencies on the order of 500Mhz or more. At various times during data processing operations in such a high speed, high bandwidth or wide bus topology, errors on the system bus may occur. Examples of such errors include a response bus error where no snooper has accepted the address; parity errors on the address/response data bus; and internal error correction code or ECC errors. It should be understood that the foregoing errors are by way of example and that other system bus errors may occur as well and invoke system bus recovery according to the present invention.

Figure 8 illustrates operation of the system of Figure 1 during system bus recovery according to the present invention. In the flow diagram of Figure 8, the following nomenclature is assigned:

REFERENCE	IDENTIFIES
H	high-frequency, wide bus - normal interconnect 12
L	low-frequency, thin bus - recovery bus of agents A1, B1, etc.
error	system errors that would normally cause machine to shut down (ECC errors, parity errors, etc., as noted above)
throttle-pls	periodic pulse that causes system to load new frequency
snp init pkt	snoop init packet on a given chip
AO pkt	address-only packet
drain H/L complete	allow address ops that are already in-flight to complete
switch system configs	switch system configuration registers from H-L settings or L-H setting via mux select
POR complete	Power-On-Reset sequence complete; meaning system is operational.

In a step 800 shown in process flow P (Figure 8), the system 8 remains in an idle state until a POR complete indication is given upon the power-on-reset sequence being completed indicating that system 8 is operational. Next a step 802 occurs where operation of the system 8 over the high frequency wide bus (H) of interconnect 12 is initiated. In the event that a snoop init packet occurs in one of the nodes, a step 804 occurs and the system begins normal operations on the wide bus H which then proceed. Should at some time during these operations of step 804 an error be detected, or in the event of receipt of a periodic pulse 'throttle-pls' occurs, causing the system 8 to load a new operating frequency, an AO or address-only packet is sent on the wide bus H.

A step 806 is initiated, to stop operations on the wide bus H and to allow a drain H operation to occur where address operations that are presently under way are allowed to complete before performance of a step 808. If desired, a programmable counter could be used in step 806 for this purpose to allow a specified time to elapse

before step 808. A suitable such time would be the time required for two address laps around the ring or loop of system 8.

During step 808, the low speed or recovery bus L is initialized, and operation of the system 8 over the recovery bus L begins. In the event a snoop init packet occurs, the configuration registers in the system 8 switch to settings for operation at a lower frequency on the low speed recovery bus from those on the wide bus H via a multiplex or mux select technique. The settings contained in the configuration registers are user defined and include settings for operation of all of the agents or nodes which are participating on the wide bus, or in some cases settings for a designated or selected fewer number of agents or nodes. In operation at the lower frequency, the number of agents or nodes selected would preferably include all agents or nodes to continue and participate. However, if the recovery time or transition time is relatively short before the high frequency bus is again operating at a new frequency, only selected agents or nodes need be indicated by the configuration registers as included on the low frequency bus. In that event the configuration registers would indicate only those selected agents or nodes.

A step 810 then occurs and the system 8 is able to continue data processing operations on the lower frequency, narrow bandwidth recovery bus L. If during step 810 a 'throttle-pls' occurs, an AO or address only packet is sent over the low speed recovery bus.

A step 812 then occurs to stop operations on the recovery bus L and to allow a drain L operation to be performed where data processing steps then underway on the recovery bus L are allowed to be completed. As was the case with step 806, a programmable counter may be used to allow a specified time to elapse during the drain L operation of step 812. An AO packet is then sent over the wide bus H of the system 8 and the process returns to step 802 where operations proceed in the manner described above.

As has been described, the present invention provides an improved non-hierarchical interconnect for a multi-node data processing system. The interconnect

architecture introduced by the present invention has an associated communication protocol having a distributed combined response mechanism that accumulates per-node partial combined responses until a complete combined response can be obtained and provided to all nodes. For both read and write communication scenarios, the combined response, in addition to conveying the snoop response of a servicing agent, indicates the node ID of the node containing the servicing agent. In this manner, read and write data can be directed from a source agent to a target agent without being propagated to other nodes unnecessarily. The present invention also introduces two mechanisms to facilitate better communication queue management: a cancel mechanism to enable remote nodes to ignore a request that can be serviced locally and a speculative cancellation mechanism that enables an agent to speculatively cancel a queue allocated to a request in response to the partial combined response for the request.

The system and method of the present invention provide an ability to dynamically increase the performance of high speed data processing systems. The operating frequency of the system 8 over the wide bus interconnect 12 may be increased as the system hardware is actually operating in real time until an optimal operating system frequency is achieved over the wide bus. This can be done without the risk of system errors causing the system 8 to completely fail. In the event system errors occur, the system and method of the present invention transfer communication from the high speed, wide bus topology to the recovery bus, as described above, where data processing continues while appropriate adjustments are made to the wide bus interconnect 12. The recovery system and method of the present invention can be used to increase system performance. The process is achievable in real time and is thus truly dynamic.

If desired, the recovery bus R need not be implemented as a wired bus. Instead, a low frequency, low speed wireless bus or virtual bus with appropriate encryption and security measures could be used as the recovery bus R in place of the wired recovery bus. Further, although the recovery technique is disclosed in the

foregoing preferred embodiment as implemented in a ring configuration of bus topology, it should be understood that the present invention could be used in hierarchical bus topologies as well.

While the invention has been particularly shown and described with reference to a preferred embodiment, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention. For example, although the present invention has been described with respect to embodiments of multi-node data processing systems, it should be understood that the interconnect architecture disclosed herein, which includes multiple uni-directional non-blocking address channels and one or more uni-directional data channels, can also be advantageously applied to single node data processing systems including multiple discrete agents (e.g., processors) and to single-chip multiprocessor data processing systems.

Although the invention has been described with reference to specific embodiments, this description is not meant to be construed in a limiting sense. Various modifications of the disclosed embodiment, as well as alternative embodiments of the invention, will become apparent to persons skilled in the art upon reference to the description of the invention. It is therefore contemplated that such modifications can be made without departing from the spirit or scope of the present invention as defined in the appended claims.